

REMOTE AUDIO DEVICE MANAGEMENT SYSTEM


INVENTORS

Qiong Liu
Don Kimber
Jonathan Foote
Chunyuan Liao
John Adcock

CERTIFICATE OF MAILING BY "EXPRESS MAIL" UNDER 37 C.F.R. §1.10

"Express Mail" mailing label number: EV 073 803 459 US
Date of Mailing: July 2, 2003

I hereby certify that this correspondence is being deposited with the United States Postal Service, utilizing the "Express Mail Post Office to Addressee" service addressed to Mail Stop PATENT APPLICATION, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 and mailed on the above Date of Mailing with the above "Express Mail" mailing label number.



Tina Galdos
Signature Date: July 2, 2003.

(Signature)

REMOTE AUDIO DEVICE MANAGEMENT SYSTEM

INVENTORS

Qiong Liu
Don Kimber
Jonathan Foote
Chunyuan Liao
John Adcock

Copyright Notice

[0001] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

Cross Reference to Related Applications

[0002] The present application is related to the following United States Patents and Patent Applications, which patents/applications are assigned to the owner of the present invention, and which patents/applications are incorporated by reference herein in their entirety:

[0003] United States Patent Application No. 10/205,739, entitled "Capturing and Producing Shared Resolution Video," filed on July 26, 2002, Attorney Docket No. FXPL-1037US0, currently pending.

Field of the Invention

[0004] The current invention relates generally to audio and video signal processing, and more particularly to acquiring audio signals and providing high quality customized audio signals to a plurality of remote users.

Background of the Invention

[0005] Remote audio and video communication over a network is increasingly popular for many applications. Through remote audio and video access, students can attend classes from their dormitories, scientists can participate in seminars held in other countries, executives can discuss critical issues without leaving their offices, and web surfers can view interesting events through webcams. As this technology develops, part of the challenge is to provide customized audio to a plurality of users.

[0006] Many audio enhancement techniques, such as beam forming and ICA (Independent Component Analysis) based blind source separation, have been developed in the past. To use these techniques in a real environment, it is critical to know spatial parameters of users' attention. For example, if the system points a high performance beam former in an incorrect direction, the desired audio may be greatly attenuated due to the high performance of the beam former. The ICA approach has similar results. If an ICA system is not configured with

information related to what a user wants to hear, the system may provide a reconstructed source signal that shields out the user's desired audio.

[0007] One common form of remote 2-way audio communication is the telephone. Telephone systems give us the opportunity to form a customized audio link with phones. To form telephone links with various collaborators, users are forced to remember large quantities of phone numbers. Although modern advanced telephones try to assist users by saving these phone numbers and corresponding collaborators' names in phone memory, going through a long list of names is still a cumbersome task. Moreover, even if a user has the number of a desired collaborator, the user does not know if the collaborator is available for a phone conversation.

[0008] Many audio pick-up systems of the prior art use far-field microphones. Far-field microphones pick up audio signals from anywhere in an environment. As audio signals come from all directions, it may pick up noise or audio signals that a user does not want to hear. Due to this property, a far-field microphone generally has worse signal-to-noise ratio than close-talking microphones. Although a far-field microphone has the drawback of a poor signal-to-noise ratio, it is still widely used for teleconference purposes because remote users may conveniently monitor the audio of an entire environment.

[0009] To overcome some of the drawbacks of far-field microphones, such as the pick-up or capture of audio signals from several sources at the same time, some researchers proposed to use the ICA approach to separate sound signals blindly for sound quality improvement. The ICA approach showed some improvement in many constraint experiments. However, this approach also raised new problems when used with far-field microphones. ICA requires more microphones than sound sources to solve the blind source separation problem. As the number of microphones increases, the computational cost becomes prohibitive for real time applications.

The ICA approach also requires its user to select proper nonlinear mappings. If these nonlinear mappings cannot match input probability density functions, the result will not be reliable.

[0010] Removing independent noises acquired by different microphones is another problem for the ICA approach. As an inverse problem, if the underlying audio mixing matrix is singular, the inverse matrix for ICA will not be stable. Besides all these problems, classical ICA approach eliminates location information of sound sources. Since the location information is eliminated, it becomes difficult for some final users to select ICA results based on location information. For example, an ideal ICA machine may separate signals from ten audio sources and provide ten channels to a user. In this case, the user must check all ten channels to select the source that the user wants to hear. This is very inconvenient for real time applications.

[0011] Besides the ICA approach, some other researchers use the beam-forming technique to enhance audio in a specific direction. Compared with the ICA approach, the beam-forming approach is more reliable and depends on sound source direction information. These properties make beam-forming better suited for teleconference applications. Although the beam-forming technique can be used for pick-up of audio signals from a specific direction, it still does not overcome many drawbacks of far-field microphones. The far-field microphone array used by a beam-forming system may still capture noises along a chosen direction. The audio “beam” formed by a microphone array is normally not very narrow. An audio “beam” wider than necessary may further increase the noise level of the audio signal. Additionally, if a beam former is not directed properly, it may attenuate the signal the user wants to hear.

[0012] FIG. 1 illustrates a typical control structure 100 of an automatic beam former control system of the prior art. Here, the control unit 140 (performed by a computer or processor) acquires environmental information 110 with sensors 120, such as microphones and video

cameras. The microphones used for the control may be the microphones used for beam-forming. A single sensor representation is illustrated to represent both audio and visual sensors to make the control structure clear. Based on the audio and visual sensory information, the control unit 140 may localize the region of interest, and point the beam former 130 to the interesting spot. In this system, the sensors and the controlled beam former must be aligned well to achieve quality audio output. This system also requires a control algorithm to accurately predict the region in which audience members are interested. Computer prediction of the region of interest is a considerable problem.

[0013] Figure 2 shows the control structure 200 of a traditional human operated audio management system. Here, the human operator 230 continuously monitors environment changes via audio and video sensors 220, and adjusts the magnification of various microphones based on environment changes. Compared to state-of-the-art automatic microphone management, a human controlled audio system is often better at selecting meaningful high quality audio signals. However, human controlled audio systems require people to continuously monitor and control audio mixers and other equipment.

[0014] What is needed is a audio device management system that enhances audio acquisition quality by using human suggestions and learning audio pick-up strategies and camera management strategies from user operations and input.

Summary of the Invention

[0015] An audio device management system (ADMS) manages remote audio devices via user selections in video links. The system enhances audio acquisition quality by receiving and processing human suggestions, forming customized two-way audio links according to user requests, and learning audio pickup strategies and camera management strategies from user operations.

[0016] The ADMS is constructed with microphones, speakers, and video cameras. The ADMS control interface for a remote user provides a multi-window GUI that provides an overview window and selection display window. With the ADMS, GUI remote users can indicate their visual attentions by selecting regions of interest in the overview window.

[0017] The ADMS provides users with more flexibility to enhance audio signals according to their needs and makes it more convenient to form customized two-way audio links without requiring users to remember a list of phone numbers. The ADMS also automatically manages available microphones for audio pickup based on microphone sound quality and the system's past experience when users monitor a structured audio environment without explicitly expressing their attentions in the video window. In these respects, the ADMS differs from fully automatic audio pickup systems, existing telephone systems, and operator controlled audio systems.

Brief Description of the Drawings

[0018] FIGURE 1 is an illustration of an automatic beam former control system of the prior art.

[0019] FIGURE 2 is an illustration of a human-operator controlled audio management system of the prior art.

[0020] FIGURE 3 is an illustration of an environment having audio and video sensors in accordance with one embodiment of the present invention.

[0021] FIGURE 4 is an illustration of a graphical user interface for providing audio and video to a user in accordance with one embodiment of the present invention.

[0022] FIGURE 5 is an illustration of a method for determining audio device selection in accordance with one embodiment of the present invention.

[0023] FIGURE 6 is an illustration of a method for providing audio based on user input in accordance with one embodiment of the present invention.

[0024] FIGURE 7 is an illustration of a method for selecting an audio source in accordance with one embodiment of the present invention.

[0025] FIGURE 8 is an illustration of a single-user controlled audio device management system in accordance with one embodiment of the present invention.

[0026] FIGURE 9 is an illustration of user selection of audio requests over a period of time in accordance with one embodiment of the present invention.

[0027] FIGURE 10 is an illustration of a cylindrical coordinate system in accordance with one embodiment of the present invention.

[0028] FIGURE 11 is an illustration of a video frame with highlighted user selections in accordance with one embodiment of the present invention.

[0029] FIGURE 12 is an illustration of a probability estimation of user selections in accordance with one embodiment of the present invention.

[0030] FIGURE 13 is an illustration of a video frame with a highlighted system selection in accordance with one embodiment of the present invention.

[0031] FIGURE 14 is an illustration of video frame with an alternative highlighted system selection in accordance with one embodiment of the present invention.

Detailed Description

[0032] Audio pickup devices used can be categorized as far-field microphones or close-talking (near-field) microphones. The audio device management system (ADMS) of one embodiment of the present invention uses both types of microphones for audio signal acquisition. Far-field microphones pick-up or capture audio signals from nearly any location in an environment. As audio signals come from multiple directions, they may also pick-up noise or audio signals that a user does not want to hear. Due to this property, a far-field microphone generally has worse signal-to-noise ratio than close-talking microphones. Although far-field microphones have this drawback of poor signal-to-noise ratio, it is still widely used for teleconferencing because it is convenient for remote users to monitor the whole environment.

[0033] To compensate for drawbacks inherent in far-field microphones, it is better to use close-talking microphones in the conference audio system. Close-talking microphones typically capture audio signals from nearby locations. Audio signals originating relatively far from this type of microphone are greatly attenuated due to the microphone design. Therefore, close-talking microphones normally achieve much higher signal-to-noise ratio than far-field microphones and are used to capture and provide high quality audio. Besides high signal-to-noise ratio, close-talking microphones can also help the system to separate a high-dimensional ICA problem into multiple low-dimensional problems, and associate location information with these low-dimensional problems. If close-talking microphones are used properly, they may also help the audio system capture less noise along a user selected direction.

[0034] Although close-talking microphones have many advantages over far-field microphones, close-talking microphones shouldn't be used to replace all far-field microphones in some circumstances for several reasons. Firstly, in a natural environment, people may sit or stand at

various locations. A small number of close-talking microphones may be not enough to acquire audio signals from all these locations. Secondly, intensively packing close-talking microphones everywhere is expensive. Finally, connecting too many microphones in an audio system may make the system too complicated. Due to these concerns, both close-talking microphone and far-field microphone are used in the ADMS construction. Similarly, various audio playback devices, such as headphones and speakers, are used in the ADMS construction.

[0035] After various devices are installed, the audio management system of the present invention may selectively amplify sound signals from various microphones according to selections relating to remote users' attentions. The physical location of a microphone is a convenient parameter for distinguishing one microphone from another. To use this control parameter, users can input the coordinates of a microphone, mark the microphone position within a geometric model, or provide some other type of input that can be used to select a microphone location. Since these approaches do not provide enough context of the audio environment, they are not a friendly interface for remote users. In one embodiment of the present invention, video windows are used as the user interface for managing the distributed microphone array. In this manner, remote users can view the visual context of an event (e.g. the location of a speaker) and manage distributed microphones according to the visual context. For example, if a user finds and selects the presenter in the visual context in the form of video, the system may activate microphones near the presenter to hear high quality audio. In one embodiment, to support this microphone array management approach, the ADMS uses hybrid cameras having a panoramic camera and a high resolution camera in the audio management system. In one embodiment, the hybrid camera may be a FlySPEC type cameras as disclosed in United States Patent Application No. 10/205,739, which is incorporated by reference in its entirety. These cameras are installed in the

same environment as microphones to ensure video signals are closely related to audio signals and microphone positions.

[0036] To illustrate the use of these ideas in a real environment, an audio management system may be discussed in the context of a conference room example. FIG. 3 illustrates a top view of a conference room 310 having sensor devices for use with an ADMS in accordance with one embodiment of the present invention. Conference room 310 includes front screen 305, podium 307, and tables 309. In the embodiment shown, close-talking microphones 320 are dispersed throughout the room on tables 309 and podium 307. In one embodiment, the close talking microphones may be GN Netcom Voice Array Microphones that work within 36 inches, or other close-field microphone combinations. In the audio system shown, many close-field microphones are located on tables 309 to capture voices and other audio near the tables 309. Far-field microphone arrays 330 can capture sound from the entire room. Camera systems 340 are placed such that remote users can watch events happening in the conference room. In one embodiment, the cameras 340 are FlySpec cameras. Headphones 350 may be placed at any location, or locations, in the room for a private discussion as discussed in more detail below. Loud speaker 360 may provide for one or more remote users to speak with those in the conference room. In another embodiment, the loud speakers allow any person, persons, or automated system to provide audio to people and audio processing equipment located in the conference room. If necessary, extending the ADMS to allow text exchange via PDA or other devices is also possible.

[0037] In one embodiment, the ADMS of the present invention may be used with a GUI or some other type of interface tool. FIG. 4 illustrates an ADMS GUI 400 in accordance with one embodiment of the present invention. The ADMS GUI 400 consists of a web browser window

410. The web browser window 410 includes an overview window 420 and a selection display window 430. The overview window may provide an image or video feed of an environment being monitored by a user. The selection display window provides a close-up image or video feed of an area of the overview window. In one embodiment wherein the video sensors include a hybrid camera such as the FlySpec camera, overview window 420 displays video content captured by the hybrid camera panoramic camera and selection display window 430 displays video content captured by the hybrid camera high resolution camera.

[0038] Using this GUI, the human operator may adjust the selection display video by providing input to select an interesting region in the overview window. Thus, a region in the overview window selected by a user generated gesture input is displayed in higher resolution in the selection display window. In one embodiment, the input may be gesture. A gesture may be received by the system of the present invention through an input device or devices such as a mouse, touch screen monitor, infra-red sensor, keyboard, or some other input device. After the interesting region is selected in some way, the region selected will be shown in the selection display window. At the same time, audio devices close to the selected region will be activated for communication. In one embodiment, the region selected by a user will be visually highlighted in the overview window in some manner, such as with a line or a circle around the selected area. For pure audio management, the selected region in the overview window is enough for the ADMS. The selection result window in the interface is to motivate the user to select her/his interested region in the upper window, and let the audio management system in the environment take control of they hybrid camera. A selection result window also helps the audio management by letting users watch more details.

[0039] In one embodiment, two modes can be configured for the interface. In the first mode, a participant or user receives one-way audio from a central location having sensors. In the embodiment illustrated in FIG. 3, the central location would be the conference room having the microphones and video cameras. When the participant selects this mode, his or her selection in the video window will be used for audio pickup. In the second mode, a remote participant or user may participate in two way audio communication with a second participant. In one embodiment, the audio communication may be with a second participant located at the central location. The second participant may be any participant at the central location. When a remote participant selects this mode, his/her selection in the video window will be used for activating both the pickup and the playback devices (e.g. a cell phone) near the selected direction.

[0040] In one embodiment, multiple users can share cameras and audio devices in the same environment. The multiple users can view the same overview window content and select their own content to be displayed in the selection result window. FIG. 5 illustrates a method 500 for implementing an ADMS control system in accordance with one embodiment of the present invention. Method 500 begins with start step 505. Next, the system determines if a user request for audio has been received in step 510. In one embodiment, the user request may be received by a user selection of a region of the overview window in ADMS GUI 400. The selection maybe input by entering window coordinates, selecting a region with a mouse, or some other means. If a user request has been received, audio is provided to the requesting user based on the user's request at step 520. Step 520 is discussed in more detail below with respect to FIG. 6. If no user request is determined to be received at step 510, then operation continues to step 530. At step 530, audio is provided to users via a rule-based system. The rule-based system is discussed in more detail below.

[0041] FIG. 6 illustrates a method 600 for providing audio to a user based on a request received from the user. Method 600 begins with start step 605. Next, an area associated with a user's selection is searched for corresponding audio devices at step 610. In one embodiment, the selection area is determined when a user selects a portion of a GUI window. The window may display a representation of some environment. The environment representation may be a video feed of some location, a still image of a location, a slide show of a series of updated images, or some abstract representation of an environment. In the GUI illustrated in FIG. 4, a user selects a portion of the overview window. In any case, different portions of the environment representation can be associated with different audio devices. The audio devices may be listed in a table or database format in a manner that associates them with specific coordinates in the GUI window. For example, in an environment representation of a conference room, wherein the window displays a speaker at a podium in the center region of the window, pixels associated with the center region of the window may be associated with output signal information regarding the microphone located at the podium. Once a selection area is received, the ADMS may search a table, database, or other source of information regarding audio devices associated with the selected area. In one embodiment, an audio device may be associated with a selected area if the audio device is configured to point, be directed to, or otherwise receive audio that originates or is otherwise associated with the selected area.

[0042] Next, the system determines if any audio devices were associated with the selected area at step 620. If audio devices are associated with the selected area, then two way communication is provided at step 630 and method 600 ends at step 660. Providing two-way communication at step 630 is discussed below with respect to FIG 7. If no audio device is found to be associated with the specific area, then operation continues to step 640 where an alternate device is selected.

The alternate device may be a device that is not specifically targeted towards the selected area but provides two way communication with the area, such as a nearby telephone. Alternatively, the alternate communication device could be a loud speaker or other device that broadcasts to the entire environment. Once the alternate audio device is selected, the alternate audio device is configured for user communication at step 650. Configuring the device for user communication includes configuring the capabilities of the device such that the user may engage in two-way audio communication with a second participant at the central location. After step 650, operation ends at step 655.

[0043] FIG. 7 illustrates a method 700 for selecting an audio device associated with a user selection in accordance with one embodiment of the present invention. Method 700 begins with start step 705. Next, the ADMS determines if more than one audio device is associated with the user selected region at step 710. If only one device is associated with the user selected region, then operation continues to step 740. If multiple devices are associated with the selected region, then operation continues to step 720. At step 720, parameters are compared to determine which of the multiple devices would be the best device. In one embodiment, parameters regarding preset security level, sound quality, and device demand may be considered. When multiple parameters are compared, each parameter may be weighted to give an overall rating for each device. In another embodiment, parameters may be compared in a specific order. In this case, subsequent compared parameters may only be compared if no difference or advantage was associated with a previously compared parameter. Once parameters associated with the audio devices are compared, the best match audio device is selected at step 730 and operation continues to step 740.

[0044] The device is activated at step 740. In one embodiment, activating a device involves providing the audio capabilities of the device to the user selecting the device. User contact information may then be provided at step 750. In one embodiment, the user contact information is provided to the audio device itself in a form that allows a connection to be made with the audio device. In another embodiment, providing contact information includes providing identification and contact information to the audio device, such that a second participant near the audio device may engage in audio communication with the first remote participant who selected the area corresponding the particular audio device. Once contact information is provided, operation of method 700 ends at step 755.

[0045] FIG. 8 illustrates a single-user controlled ADMS 800 in accordance with one embodiment of the present invention. ADMS 800 includes environment 810, sensors 820, computer 830, human 840, coordinator 850, and audio server 860.

[0046] In this system, both the human operator (i.e., the system user) and the automatic control unit can access data from sensors. In one embodiment of the present invention, the sensors may include panoramic cameras, microphones, and other video and audio sensing devices. With this system, the user and the automatic control unit can make separate decisions based on environmental information. In one embodiment, the decisions by the user and automatic control unit may be different. To resolve conflicts, the human decision and the control unit decision are sent to a coordinator unit before the decision is sent to the audio server. In a preferred embodiment, the human choice is considered more desirable and meaningful than the automatic selection. In this case, a human decision in conflict with an automatic unit decision overrides the automatic unit decision inside the coordinator. In another embodiment, each of the user and automatically selected regions are associated with a weight. Factors in determining the weight of

each selection may include signal-to-noise ratio in the audio associated with each selection, reliability of the selection, the distortion of the video content associated with each selection, and other factors. In this embodiment, the coordinator will select the selection associated with the highest weight and provide the audio corresponding to the weighted selection to the user. In an embodiment where no user selection is made within a certain time period, the weight of the user selection is reduced such that the automatic selection is given a higher weight.

[0047] In ADMS 800, the user monitors the microphone array management process instead of operating the audio server continuously. To ensure audio selection quality, the human operator only needs to adjust the system when the automatic system misses the direction of interest. Thus, the system is fully automatic when no human operator provides controlling input. For an automatic system, which may miss the correct direction for audio enhancement, a human operator can drastically decrease the miss rate. Compared with a manual microphone array management system, this system can substantially reduce the human operator effort required. ADMS 800 allows users to make the tradeoff between operator effort and audio quality.

[0048] With the control structure setup illustrated in FIG. 8, audio management is performed by maximizing the audio quality in user-selected directions. As multiple users access the ADMS simultaneously, the ADMS generates multiple optimal audio signal streams for various users according to their respective requests. In one embodiment, the ADMS of the present invention measures audio quality with signal-to-noise ratio. Assume i is the index of microphones, s_i is the pure signal picked by microphone i , n_i is the noise picked by microphone i , (x_i, y_i) is the coordinates of microphone i 's image in the video window, and R_u is the region related to a user u 's selection in the video window. A simple microphone selection strategy for user u can be defined with

$$i_u = \arg \max_{(x_i, y_i) \in R_u} \left(\frac{s_i}{n_i} \right) \quad (1)$$

[0049] Thus, equation (1) selects the microphone or other audio signal capturing device which has the best signal-to-noise ratio (SNR) in the user-selected region or direction. Thus, the microphone may be located in the area corresponding to the region selected by the user or be directed to capture audio signals present in the region selected by the user. In this equation, the definition of R_u may be defined in a static or dynamic way. The simplest definition of R_u is the user-selected region. For a fixed close-talking microphone, such as microphone 320 shown in FIG. 3, the coordinates of the microphone in the window are fixed. For a far-field microphone array near to a video camera, such as microphone 330 shown in FIG. 3, its coordinates may be anywhere in the corresponding video window supported by camera 340 in FIG. 3. A far-field microphone that is not near a camera is considered to be a microphone that can be moved anywhere. Therefore, the optimization in eq. (1) takes both far-field microphones and near-field microphones into account. In another embodiment, a more sophisticated definition of R_u may be the smallest region that includes k microphones around the selected region center. When a user does not make any selection, the system can pick the microphone for this user according to

$$i_u = \arg \max_{(x_i, y_i) \in \{R_{u1}, R_{u2}, \dots, R_{uM}\}} \left(\frac{s_i}{n_i} \right) \quad (2)$$

[0050] This is the best channel within all users' selections $\{R_{u1}, R_{u2}, \dots, R_{uM}\}$. When no user gives any suggestion to the microphone management system, the selection can be over all microphones. This selection can be described with

$$i_u = \arg \max \left(\frac{s_i}{n_i} \right) \quad (3)$$

[0051] The audio system of the present invention may use other audio device selection techniques, such as ICA and beam forming. For example, K number of microphones can be used near the selected region to perform ICA. The K signals can also be shifted according to their phases, and can be added together to reduce unwanted noises. All outputs generated by ICA and beam forming may be compared with the original K signals. Regardless of the method used, the determination for final output may still be based on SNR.

[0052] From eq. (1)-(3), it is assumed that signal and noise are known for each microphone. In an embodiment wherein signal and noise are not known for a microphone, a threshold for the microphone can be set. In one embodiment, the threshold may be set according to experiment, wherein acquired data is considered noise if the data is below the threshold. In this way, the system may estimate the noise spectrum $n_i(f)$ when no event is going on or minimal audio signals are being captured by microphones and other devices. When the microphone acquires data $a_i(f)$ that is higher than the threshold, the signal spectrum $s_i(f)$ may be estimated with

$$s_i(f) = \begin{cases} 0 & \text{if } [a_i(f) - n_i(f)] < 0 \\ a_i(f) - n_i(f) & \text{if } [a_i(f) - n_i(f)] \geq 0 \end{cases} \quad (4)$$

[0053] When noise estimations are available for every microphone, the processing steps are similar to that for estimating noises and signals of all ICA outputs and beam-forming outputs. In one embodiment, the ADMS of the present invention may learn from user selections over time. User operations provide the system precious data about users' preferences. The data may be used by ADO to improve itself gradually. The ADMS may employ a learning system run in parallel with the automatic control unit, so it can learn audio pickup strategies from human user operations. In one embodiment, a_1, a_2, \dots, a_R represent measurements from environmental

sensors, and (x,y) on the captured main image correspond to a position of interest. In one embodiment, the main image may be a panoramic image. Then, the destination position (X,Y) for the audio pickup can be estimated with:

$$\begin{aligned}
 (X,Y) &= \arg \max_{(x,y)} \{ p[(x,y) | (a_1, a_2, \dots, a_R)] \} \\
 &= \arg \max_{(x,y)} \left\{ \frac{p[(a_1, a_2, \dots, a_R | (x,y)] \cdot p(x,y)}{p(a_1, a_2, \dots, a_R)} \right\} \\
 &= \arg \max_{(x,y)} \{ p[(a_1, a_2, \dots, a_R) | (x,y)] \cdot p(x,y) \}
 \end{aligned} \tag{5}$$

[0054] Assuming a_1, a_2, \dots, a_R are conditionally independent, the camera position can be estimated with:

$$\begin{aligned}
 (X,Y) &= \arg \max_{(x,y)} \{ p[(x,y) | (a_1, a_2, \dots, a_R)] \} \\
 &= \arg \max_{(x,y)} \{ p[a_1 | (x,y)] \cdot p[a_2 | (x,y)] \cdots p[a_R | (x,y)] \cdot p(x,y) \}
 \end{aligned} \tag{6}$$

[0055] The probabilities in eq. (6) can be estimated online. For example, FIG. 9 shows the users' selections during an extended period of a meeting for which the probability $p(x,y)$ is being estimated. A typical image recorded during the meeting is used as the background to illustrate the spatial arrangement of a meeting room. In this figure, users' selections are marked with boxes. Many boxes in the image form a cloud of users' selections in the central portion of the image, where the presenter and a wall-sized presentation display are located. Based on this selection cloud, it is straightforward to estimate $p(x,y)$.

[0056] Using progressive learning enables the system of the present invention to better adapt to environmental changes. In some cases, some sensors may become less reliable. For example, desks being moved may block the sound path of a microphone array. To adapt to these changes,

a mechanism can learn how informative each sensor is. Assume (U,V) is the position of interest estimated by a sensor (a camera, microphone array, or other audio capture device) and (X,Y) is the camera position decided by users. How informative the sensor is can be evaluated through online estimation as follows:

$$I[(U,V),(X,Y)] = \int_{(U,V),(X,Y)} p[(U,V),(X,Y)] \cdot \log \frac{p[(U,V),(X,Y)]}{p(U,V) \cdot p(X,Y)} \quad (7)$$

[0057] Evaluation of eq. (7) gives mutual information between (U,V) and (X,Y) . The higher the value, the more important the sensor is to the automatic control. When a sensor is broken, disabled, or yields poor information for any reason, the mutual information between the sensor and the human selection will decrease to a very small value, and the sensor will be ignored by the control software. This is helpful in allocating computational power to useful sensors. With similar techniques, the system can disable the rule-based automatic control system when the learning system can operate the camera better.

[0058] The signal quality of the captured audio signal can be processed and measured in numerous ways. In one embodiment, the signal quality of the audio signal may be improved by attempting to reduce the distortion of the audio signal captured.

[0059] Conceptually, the ideal signal received at a given point may be represented with $f(\phi,\theta,t)$, where ϕ and θ are spatial angles used to identify the direction of a coming signal and t is the time. For derivations in later applications, a cylindrical coordinate system 1000 illustrated in FIG. 10 may be used to describe the signal. In Figure 10, a line passing through the origin and a point on a cylindrical surface is used to define the signal direction. The point on the cylindrical surface

has coordinates (x,y) , where x is the arc length between $(x=0, y=0)$ and the point's projection on $y=0$, and y is the height of the point from the plane $y=0$. With this coordinate system, the ideal signal is represented with $f(x,y,t)$. In one embodiment, a signal acquisition system may capture an approximation $\hat{f}(x,y,t)$ of the ideal signal $f(x,y,t)$ due to the limitation of sensors. The sensor control strategy in one embodiment is to maximize the quality of the acquired signal $\hat{f}(x,y,t)$.

[0060] The information loss of representing f with \hat{f} may be defined with

$$D[\hat{f}, f] = \sum_i p(R_i, t | O) \iiint_{R_i, T} |\hat{f}(x, y, t) - f(x, y, t)|^2 dx dy dt, \quad (8)$$

where $\{R_i\}$ is a set of non-overlapping small regions, T is a short time period, and $p(R_i, t | O)$ is the probability of requesting details in the direction of region- R_i details (conditioned on environmental observation O).

[0061] This probability may be obtained directly based on users' requests. Suppose there are $n_i(t)$ requests to view region R_i during the time period from t to $t+T$ when the observation O is presented, and p and O do not change much during this period, then $p(R_i, t | O)$ may be estimated as

$$p(R_i, t | O) = \frac{n_i(t)}{\sum_i n_i(t)}. \quad (9)$$

$\iiint_{R_i, T} |\hat{f}(x, y, t) - f(x, y, t)|^2 dx dy dt$ is easier to estimate in the frequency domain. If ω_x and ω_y represent spatial frequencies corresponding to x and y respectively, and ω_t is the temporal frequency, the distortion may be estimated with

$$\begin{aligned}
& \iiint_{R_i, T} \left| \hat{f}(x, y, t) - f(x, y, t) \right|^2 dx dy dt \\
&= \iiint_{R_i, T} \left| \hat{F}(\omega_x, \omega_y, \omega_t) - F(\omega_x, \omega_y, \omega_t) \right|^2 d\omega_x d\omega_y d\omega_t
\end{aligned} \tag{10}$$

[0062] The accomplishment of acquiring a high quality signal is equivalent to reducing $D[\hat{f}, f]$.

Assume $\hat{f}(x, y, t)$ is a band limited representation of $f(x, y, t)$. Reducing $D[\hat{f}, f]$ may be achieved by moving steerable sensors to adjust cutoff frequencies of $\hat{f}(x, y, t)$ in various regions $\{R_i\}$.

Assume the region i of $\hat{f}(x, y, t)$ has spatial cutoff frequencies $a_{x,i}(t)$, $a_{y,i}(t)$, and temporal cutoff frequency $a_{t,i}(t)$. The estimation of $\iiint_{R_i, T} \left| \hat{f}(x, y, t) - f(x, y, t) \right|^2 dx dy dt$ may then be simplified to

$$\begin{aligned}
& \iiint_{R_i, T} \left| \hat{f}(x, y, t) - f(x, y, t) \right|^2 dx dy dt \\
&= \iiint_{\substack{R_i, T \\ \omega_x > a_{x,i}(t) \\ \omega_y > a_{y,i}(t) \\ \omega_t > a_{t,i}(t)}} \left| \hat{F}(\omega_x, \omega_y, \omega_t) - F(\omega_x, \omega_y, \omega_t) \right|^2 d\omega_x d\omega_y d\omega_t
\end{aligned} \tag{11}$$

[0063] In this embodiment, the optimal sensor control strategy is to move high-resolution (i.e. in space and time) sensors to certain locations at certain time periods so that the overall distortion $D[\hat{f}, f]$ is minimized.

[0064] Equations (8)-(11) described a way to compute the distortion when participants' requests were available. When participants' requests are not available, the estimation of $p(R_i, t | O)$ may become a problem. This may be overcome by using the system's past experience of users' requests. Specifically, assuming that the probability of selecting a region does not depend on time t , the probability may be estimated as:

$$p(R_i, t | O) = p(R_i | O) = \frac{p(O | R_i) \cdot p(R_i)}{p(O)} . \tag{12}$$

[0065] O can be considered an observation space of \hat{f} . By using a low dimensional observation space, it is easier to estimate $p(R_i, t | O)$ with limited data. With this probability estimation, the

system may automate the signal acquisition process when remote users don't, won't, or cannot control the system.

[0066] The equations (8)-(12) can be directly used for active sensor management. For better understanding of the present invention according to one embodiment, a conference room camera control example can be used to demonstrate the sensor management method of this embodiment of the present invention. A panoramic camera was used to record 10 presentations in our corporate conference room and 14 users were asked to select interesting regions on a few uniformly distributed video frames, using the interface shown in Figure 4. Figure 11 shows a typical video frame and corresponding selections highlighted with boxes. Figure 12 shows the probability estimation based on these selections. In Figure 12, lighter color corresponds to higher probability value and darker color corresponds to lower value.

[0067] To compute the distortion defined with eq.(8), the system needs the result from eq.(11). Since it is impossible to get complete information of $F(\omega_x, \omega_y, \omega_t)$, the system needs proper mathematical models to estimate the result. According to Dong and Atick, "Statistics of Natural Time Varying Images", Network:Computation in Neural Systems, vol. 6(3), pp.345-358, 1995, if a system captures object movements from distance zero to infinity, $F(\omega_x, \omega_y, \omega_t)$ statistically falls with temporal frequency, ω_t , and rotational spatial frequency, ω_{xy} , according to

$$|F(\omega_{xy}, \omega_t)|^2 = \frac{A}{\omega_{xy}^{1.3} \cdot \omega_t^2}, \quad (13)$$

where A is a positive value related to the image energy.

[0068] In one embodiment, b_{xy} and b_t can be denoted as the spatial and temporal cutoff frequencies of the panoramic camera and a_{xy} and a_t as the spatial and temporal cutoff frequencies of a PTZ camera. Let

$$\begin{aligned}
E_{xy} &= \int_1^b \int_1^{b_y} |F(\omega_{xy}, \omega_t)|^2 d\omega_{xy} d\omega_t \\
E_{xy} &= \int_1^{b_y} |F(\omega_{xy}, 0)|^2 d\omega_{xy} \\
E_t &= \int_1^b |F(0, \omega_t)|^2 d\omega_t
\end{aligned} \tag{14}$$

[0069] If the system uses the PTZ camera instead of the panoramic camera to capture region R_i , the video distortion reduction achieved by this may be estimated with

$$\begin{aligned}
D_{G,i} &= \left[\frac{(a_{xy}^{0.3} - 1) \cdot (a_t - 1) \cdot b_{xy}^{0.3} \cdot b_t}{a_{xy}^{0.3} \cdot a_t \cdot (b_{xy}^{0.3} - 1)(b_t - 1)} - 1 \right] \cdot E_{xy,t} \\
&+ \left[\frac{(a_{xy}^{1.3} - 1) \cdot b_{xy}^{1.3}}{a_{xy}^{1.3} \cdot (b_{xy}^{1.3} - 1)} - 1 \right] \cdot E_{xy,i} + \left[\frac{(a_t - 1) \cdot b_t}{a_t \cdot (b_t - 1)} - 1 \right] \cdot E_{t,i}
\end{aligned} \tag{15}$$

[0070] Coordinates (X, Y, Z) , corresponding to sensor features pan/tilt/zoom, can be associated with as the best pose of the camera or sensor. With eq.(8) and eq.(15), (X, Y, Z) can be estimated with

$$(X, Y, Z) = \arg \max_{(x, y, z)} [p(R_i, t | O) \cdot D_{G,i}] \tag{16}$$

[0071] In the experiment discussed above, the panoramic camera has 1200x480 resolution, and the PTZ camera has 640x480 resolution. Compared with the panoramic camera, the PTZ camera can achieve up to 10 times higher spatial sampling rate by performing optical zoom in practice. The camera frame rate varies over time depending on the number of users and the network traffic. The frame rate of the panoramic camera was assumed to be 1 frame/sec and the frame rate of the PTZ camera is assumed to be 5 frames/sec. With the above optimization procedure and users' suggestions shown in Figure 11, the system selects the rectangular box in Figure 13 as the view of the PTZ camera.

[0072] When users' selections are not available to the system, the system has to estimate the probability term (i.e. predicts users' selections) according to eq.(13). Due to the imperfection of the probability estimation, the distortion estimation without users' inputs is a little bit different

from the distortion estimation with users' inputs. This estimation difference leads the system to a different PTZ camera view suggestion shown in Figure 14. By visually inspecting automatic selections over a long video sequence, these automatic PTZ view selections are very close to those PTZ view selections estimated with users' suggestions. If we replace the panoramic camera and the PTZ camera in this experiment with a low spatial resolution microphone array and a steer-able unidirectional microphone, the proposed control strategy can be used to control the steer-able microphone as we use it to control the PTZ camera.

[0073] Other features, aspects and objects of the invention can be obtained from a review of the figures and the claims. It is to be understood that other embodiments of the invention can be developed and fall within the spirit and scope of the invention and claims.

[0074] The foregoing description of preferred embodiments of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obviously, many modifications and variations will be apparent to the practitioner skilled in the art. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalence.

[0075] In addition to an embodiment consisting of specifically designed integrated circuits or other electronics, the present invention may be conveniently implemented using a conventional general purpose or a specialized digital computer or microprocessor programmed according to the teachings of the present disclosure, as will be apparent to those skilled in the computer art.

[0076] Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art. The invention may also be implemented by the preparation of application specific integrated circuits or by interconnecting an appropriate network of conventional component circuits, as will be readily apparent to those skilled in the art.

[0077] The present invention includes a computer program product which is a storage medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

[0078] Stored on any one of the computer readable medium (media), the present invention includes software for controlling both the hardware of the general purpose/specialized computer or microprocessor, and for enabling the computer or microprocessor to interact with a human user or other mechanism utilizing the results of the present invention. Such software may include, but is not limited to, device drivers, operating systems, and user applications.

[0079] Included in the programming (software) of the general/specialized computer or microprocessor are software modules for implementing the teachings of the present invention, including, but not limited to, remotely managing audio devices.